



HAL
open science

Noise-Tolerant Active Preference Learning for Multicriteria Choice Problems

Margot Herin, Patrice Perny, Nataliya Sokolovska

► **To cite this version:**

Margot Herin, Patrice Perny, Nataliya Sokolovska. Noise-Tolerant Active Preference Learning for Multicriteria Choice Problems. *Algorithmic Decision Theory (ADT 2024)*, Oct 2024, New Brunswick, NJ, United States. pp.191–206, 10.1007/978-3-031-73903-3_13 . hal-04841741

HAL Id: hal-04841741

<https://hal.science/hal-04841741v1>

Submitted on 16 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

NOISE-TOLERANT ACTIVE PREFERENCE LEARNING FOR MULTICRITERIA CHOICE PROBLEMS

Margot Herin, Patrice Perny
LIP6
Sorbonne University,
Paris
{margot.herin, patrice.perny}@lip6.fr

Nataliya Sokolovska
LCQB
Sorbonne University,
Paris
nataliya.sokolovska@sorbonne-universite.fr

ABSTRACT

To make a choice in the presence of multiple criteria, we generally use an aggregation function which determines, for each alternative, the balance of its strengths and weaknesses and its overall evaluation. The aggregation function uses weights to adapt the model to the decision-maker's value system, by specifying the importance of the criteria and possibly their interactions. In this paper, we propose a noise-tolerant active learning method for these parameters, which not only effectively reduces the indeterminacy of the weights to identify an optimal or near-optimal decision among a given set of alternatives, but also simultaneously determines a predictive model of preferences capable of making relevant choices for the decision-maker on new instances. These outcomes are achieved by leveraging a general disagreement-based active learning approach that is theoretically guaranteed to be tolerant to noisy answers. The proposed method applies to various weighted aggregation functions, linear or not, classically used in decision theory.

1 Introduction

In multicriteria choice problems, it is commonly accepted that the exploration of admissible trade-offs should be restricted to Pareto-optimal solutions, i.e. solutions that cannot be improved on one criterion without having to be degraded on another. These solutions are, however, potentially very numerous, and it is necessary to collect additional preference information to define how the evaluations from the different criteria combine to define the overall preference. Decision theory provides numerous mathematical models to account for agents' preferences, in particular scalarizing functions which associate with any vector of partial evaluations a global evaluation defining the overall utility of the solution for the decision-maker (DM). The multicriteria choice problem can then be reformulated as a problem of maximizing the scalarizing function over all feasible performance vectors, i.e., the evaluation vectors associated with the alternatives in the problem under consideration.

In order to explore Pareto-optimal solutions and identify the optimal compromise for the DM, the scalarizing function f must be monotonically increasing with respect to Pareto dominance. That is, if a solution x is at least as good as x' on all criteria, we must have $f(x) \geq f(x')$. A wide range of such aggregation functions is available, from the simplest such as the weighted sum, to more sophisticated and expressive models such as the multilinear model and non-additive integrals (Choquet or Sugeno integrals), which can be used to model interactions between criteria, without forgetting weighted norms, which have proved effective for interactively exploring the set of non-dominated solutions in the Pareto sense. The weights used in these aggregators specify the relative importance of criteria and/or sometimes also the positive or negative interactions between criteria. They enable us to adapt a generic model to the DM's particular value system. We therefore optimize a function f_w parameterized by the weighting vector w . The phase of eliciting preferences and learning the w weighting vector is absolutely crucial, as it completely determines the nature of the compromise that will be found by optimizing f_w and the recommendation that will follow.

Various approaches have been proposed for specifying the parameters of a decision model in different contexts and thus determining an optimal choice. Here, we distinguish three types of approaches.

The local and interactive judgment approach: we choose an initial vector of parameters w , calculate an optimal solution for f_w , then let w evolve according to user feedback until we arrive at a solution that satisfies the DM. This approach, widely used in interactive multicriteria optimization [18, 19], allows a user-driven exploration of the Pareto set, alternating phases of calculation of the current optimal solution and phases of dialogue with the user. It may require numerous interactions, and the quality of the solution chosen at the end of the process is only validated by the decision-maker's instant sense of satisfaction.

Incremental preference elicitation: a first approach to incremental elicitation consists in progressively reducing the space of admissible parameters. Iteratively, a preference query is chosen, the answer to which induces a new constraint on the parameter space. The set of parameters compatible with the constraints induced by the preference judgments expressed is progressively reduced until the point where an alternative proves optimal for all remaining parameters (necessarily optimal solution). This approach is introduced in the ISMAUT method [21]. A principle of active question selection is often used, based on the minimization of maximum regret, to choose the most informative question [20, 5, 3] and derive a robust recommendation. Another approach, more tolerant to noisy responses, is to manage a probability distribution (or other uncertainty model [1]) over the parameter space and revise it according to the answers to questions, to choose a decision having the maximum expected value [6] or minimizing the expectation of regret [4]. These methods are question-saving, as they direct the questionnaire towards the resolution of a particular instance. On the other hand, they do not produce a learned model and are generally not sufficient to solve a choice problem involving a new set of alternatives.

Complete learning of the decision model: we use a base of preference examples and perform a regression (on the values or the order induced by the values) to determine the parameter w that best fits the example base [11, 16]. To determine the parameters of the aggregation function accurately and reliably, the model must be trained on a large base of examples, and requires a much larger number of preference queries than in incremental approaches. On the other hand, the learned model can be reused to treat a new choice problem with the same decision-maker, on new alternatives.

In this paper we propose a hybrid active learning approach that combines the objectives of the last two items, namely to quickly identify the optimal choice on the instance to be solved while providing a model capable of explaining the decision-maker's preferences and predicting his choices or formulating recommendations adapted to his preferences on new instances of choice problems.

Another aim of the work presented here is to improve the elicitation method's tolerance to DM's noisy answers. Indeed, the aim of minimizing the elicitation effort and the number of questions asked or examples used to determine an optimal alternative often leads to taking each answer as valid information, likely to definitively constrain the space of admissible parameters, without any further possibility of questioning or checking. This approach, which aims for questionnaire efficiency, is of course rather risky, as it omits any validation operation through partial redundancy of questions, nor any compromise between partially contradictory answers within the framework of a given decision model. The pitfalls of incremental elicitation by progressive and definitive reduction of possible parameters are well illustrated by the following example.

Example 1 Consider a set $X = \{a^0, \dots, a^q\}$ of $q + 1$ alternatives evaluated on two criteria and represented by performance vectors $a^i = (i, q - i)$ for $i = 0, \dots, q$. Suppose the DM has expressed a first preference $a^r \succ a^t$ for two indices $r, t \in \{0, \dots, q\}$ such that $r > t$. Suppose we want to learn the weights of a weighted sum model of the form $f_w(x) = wx_1 + (1 - w)x_2$ for an unknown parameter $w \in [0, 1]$. The preference $a^r \succ a^t$ implies $wa_1^r + (1 - w)a_2^r > wa_1^t + (1 - w)a_2^t$ and therefore $wr + (1 - w)(q - r) > wt + (1 - w)(q - t)$, or equivalently $w(r - t) > (1 - w)(r - t)$, hence $w > 1 - w$ and thus $w > 1/2$. Under this constraint, it's easy to see that $f_w(a^q) > f_w(a^i)$ for all $i < q$. We indeed have $f_w(a^q) - f_w(a^i) = wq - (wi + (1 - w)(q - i)) = (2w - 1)(q - i) > 0$ since $w > 1/2$ and $q > i$. Hence a^q is necessarily an optimal solution in X . Note, however, that if the decision-maker was mistaken in the first answer ($a^t \succ a^r$ being the actual preference), then the same reasoning would have led to the choice of a^0 . In this case, the recommendation a^q is in fact the worst possible recommendation given the actual DM's preferences.

Although this example is a bit of a caricature, it does illustrate that a concern for efficiency in the active choice of a question to ask can lead to choices that are not robust to noisy responses. In this paper, we will propose a non-Bayesian approach to active learning of decision-maker preferences, which is more robust to noisy responses than usual methods based on regret minimization and enables us to identify or approximate a necessary winner in a given set of alternatives, as well as to build an explanatory model of decision-maker preferences. The paper is organized as follows. Section 2 introduces background and notations. In Section 3, we introduce an algorithm for active preference learning in a noisy setting, and in Section 4, we demonstrate its benefits on synthetic preference data. Concluding remarks close the paper.

2 Background and Notations

Multicriteria decision-making problems are characterized by alternatives evaluated with respect to n dimensions representing various points of view (criteria evaluation or individual opinions) possibly conflicting with each other. In the sequel, $N = \{1, \dots, n\}$ denotes a set of criteria. Any alternative $x \in X$ is represented by an evaluation vector $x = (x_1, \dots, x_n)$ where $x_i \in [0, 1]$ represents the value of x with respect to criterion i for $i = 1, \dots, n$. Thus, the set $\mathcal{X} = [0, 1]^n$ represents the criteria space. Let us consider a choice problem over a finite set $X \subseteq \mathcal{X}$ representing all feasible evaluation vectors. The subset of Pareto optimal alternatives in X is denoted by X_P .

Let $f_w : \mathcal{X} \rightarrow \mathbb{R}$ be the scalarizing function used to represent the DM's preference \succ , i.e., :

$$x \succ x' \iff f_w(x) > f_w(x'), \quad \forall x, x' \in X \quad (1)$$

and let W be the associated definition domain of parameter w . A basic example is the weighted sum used in Example 1 (formally defined as $f_w(x) = \sum_{i=1}^n w_i x_i$ with $W = \{w \in [0, 1]^n \mid \sum_{i=1}^n w_i = 1\}$). Beyond its simplicity and intuitive appeal, this linear aggregation function may suffer from well-known limits in view of exploring Pareto-optimal tradeoffs when the set of alternatives is discrete (or more generally non-convex). Some Pareto-optimal tradeoffs are indeed impossible to obtain by maximizing a weighted sum (e.g, solutions that are interior points of the convex hull of feasible points). Such solutions are referred to as unsupported Pareto-optimal solutions. By limiting ourselves to the supported solutions we deprive the decision-maker of the possibility of examining various interesting trade-offs that could possibly better correspond to its value system. We give below two major examples of more sophisticated scalarizing functions able to cover a wider subset of Pareto-optimal tradeoffs.

The *Choquet integral* [13], used in multicriteria decision-making to model preferences in the presence of interacting criteria, employs a weighting system $w : 2^N \rightarrow \mathbb{R}$ that attaches a weight $w(S)$ to any possible set of criteria $S \subseteq N$, to aggregate the evaluation vector's values as follows:

$$f_w(x) = \sum_{i=1}^n [w(X_{(i)}) - w(X_{(i+1)})] x_{(i)} \quad (2)$$

where (\cdot) is any permutation of N such that $x_{(i-1)} \leq x_{(i)}$, $i = 1, \dots, n$, $X_{(i)} = \{(i), \dots, (n)\}$ and $X_{(n+1)} = \emptyset$. The set function w is called a capacity and is supposed to be normalized, i.e., $w(\emptyset) = 0$ and $w(N) = 1$. Also, to guarantee the monotonicity of f_w w.r.t. Pareto dominance, the capacity is supposed to be monotonic w.r.t. set inclusion, i.e., for any $A \subseteq B$, $w(A) \leq w(B)$ and thus $W = \{w : 2^N \rightarrow [0, 1] \mid w(A) \leq w(B) \text{ for all } A, B \text{ such that } A \subseteq B\}$. An alternative representation of any capacity w is given by its Möbius transform m_w defined as follows:

$$m_w(S) = \sum_{T \subseteq S} (-1)^{|S \setminus T|} w(T) \quad \text{with} \quad w(S) = \sum_{T \subseteq S} m_w(T) \quad (3)$$

The values $m_w(S)$ are called Möbius masses. We remark that we necessarily have $\sum_{T \subseteq N} m_w(T) = 1$, since $w(N) = 1$. Let us also recall that a Choquet integral admits a simple reformulation from m_w [7]:

$$f_w(x) = \sum_{S \subseteq N} m_w(S) \min_{i \in S} \{x_i\} \quad (4)$$

This formulation can be further simplified when considering k -additive capacities, i.e., capacities whose Möbius masses are zero for subsets larger than k and non-zero for at least one subset of size k . Whenever w is 1-additive, which is equivalent to $w(A) = \sum_{i \in A} w(\{i\})$, $A \subseteq N$ according to Equation 3, the Choquet integral boils down to a weighted sum. Whenever w is 2 additive, the weighted sum is augmented with a linear combination of pairwise minimum of type $\min\{x_i, x_j\}$ allowing the representation of positive or negative synergies for every pair of criteria. Obviously, Choquet integrals used with k -additive capacities ($2 \leq k \leq n$) involve larger criteria interactions as k increases, with enhanced descriptive possibilities. In particular, their maximization can lead to unsupported solutions.

Despite its enhanced descriptive power compared to the weighted sum, some non-supported Pareto-optimal solutions may not be accessible by maximizing a Choquet integral. In contexts where no prior information about the preference system of the DM is available, it may be the case that any Pareto-optimal solution is of possible interest and must be accessible by the scalarizing function. In this case, the standard approach is to minimize a weighted Chebyshev norm measuring the distance to ideal point [22]. This is equivalent to maximize the following scalarizing function:

$$f_w(x) = - \max_{i \in N} \{w_i |x_i - u_i|\} \quad (5)$$

where $u = (u_1, \dots, u_n)$ is the ideal point, defined by $u_i = \max_{x \in X_P} \{x_i\}$ and the associated weight definition domain is $W = \{w \in [0, 1]^n \mid \sum_{i=1}^n w_i = 1\}$.

The relative ability of the three families of scalarizing functions mentioned above to describe observed choices is illustrated on the simple instance of a bi-objective problem involving 12 Pareto-optimal solutions depicted in Figure 1. We can see the supported solutions (accessible by maximization of a weighted sum) (blue round) that lie on the upper part of the convex hull of the alternative set, then those unsupported solutions that are accessible by maximization of a Choquet integral (green diamond), and finally those unsupported solutions that are only accessible by maximization of a Chebyshev norm (red square).

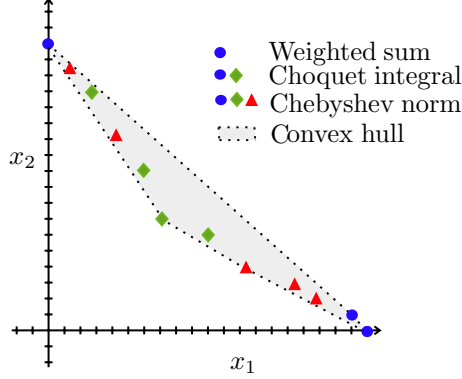


Figure 1: Accessibility of various Pareto-optimal tradeoffs

3 Noise-tolerant Active Preference Learning

Disagreement-based active learning [15, 9] refers to a branch of theoretically grounded active learning algorithms for binary classification, that share with incremental preference elicitation methods the common objective of reducing the space of admissible models as fast as possible (in terms of the number of asked questions). In the following, we establish a link between both worlds and propose a disagreement-based active preference learning for choice problems, which we illustrate on the toy case of Example 1.

3.1 Disagreement-based Active Learning

In binary classification, the input consists of data points z^i belonging to an input space \mathcal{Z} and their labels $y^i \in \mathcal{Y} = \{-1, +1\}$ distributed according to a joint distribution over $\mathcal{Z} \times \mathcal{Y}$ denoted by $\mathbb{P}_{\mathcal{Z}\mathcal{Y}}$. In this setting, based on data samples $(z^i, y^i)_{i=1}^m$, the goal is to find the classifier h^* in a hypothesis class \mathcal{H} that allows to best predict the target value of a point z , i.e., that minimizes the expectation of making a bad prediction. More formally, $h^* = \arg \min_{h \in \mathcal{H}} \ell(h)$ where $\ell(h) = \mathbb{E}_{(z,y) \sim \mathbb{P}_{\mathcal{Z}\mathcal{Y}}} [\mathbb{1}_{\{h(z) \neq y\}}] = \mathbb{P}_{(z,y) \sim \mathbb{P}_{\mathcal{Z}\mathcal{Y}}} (h(z) \neq y)$, for any $h \in \mathcal{H}$.

In this setting, the general idea of disagreement-based active learning is embodied by the CAL algorithm [8] (called after the authors' names Cohn, Atla and Ladner). Starting with \mathcal{H} as the space of admissible models (also called the version space), CAL iteratively proceeds a sequence of unlabeled points $(z^i)_{i=1}^m$, and at each iteration asks for the label y^i of the point z^i if and only if there exist two classifiers h_1, h_2 in the current space of admissible model that disagree on z^i , i.e., such that $h_1(z^i) \neq h_2(z^i)$. The portion of the input space \mathcal{Z} in which this condition holds defines the disagreement region. If asked, the newly obtained label y^i provides the additional constraint $h(z^i) = y^i$ on the set of admissible models, which now excludes the models classifying z^i differently. By doing so, the algorithm asks for a label if and only if the new constraint $h(z^i) = y^i$ surely reduces the space of admissible models, allowing to narrow down the version space around h^* with minimal labeling effort.

Obviously, this algorithm works to identify h^* under the separability hypothesis, i.e., $h^*(z^i) = y^i, i = 1, \dots, m$. In the more realistic noisy case where $\ell(h^*) > 0$ (referred to as the agnostic setting), the hard constraints $h(z^i) = y^i$ will eventually exclude h^* from the set of admissible models. Extensions of the CAL algorithm in the noisy case [2, 10] bypass this issue by defining the set of admissible models as the set of models that proved to yield small errors on the learning examples $(z^i, y^i)_{i=1}^m$. Among them, the DMH algorithm [10] (also called after the authors' names Dasgupta, Hsu, and Monteleoni), which relies on supervised learning sub-tasks, provides a simple way of cautiously excluding models associated with significantly high errors. Here we propose to exploit and extend it to the multicriteria choice problem under noisy answers.

In the multicriteria choice problem setting, the determination of the weight vector w in the preference model (1) from pairwise preference examples $x^i \succ x^{i'}$, $x^i, x^{i'} \in X_P$ can be indeed formulated as a binary classification problem where $z^i = (x^i, x^{i'}) \in \mathcal{Z} = X_P^2$ and $y^i = 1$ if $x^i \succ x^{i'}$ and $y^i = -1$ otherwise ($x^i \succsim x^{i'}$). In this case, the hypothesis class can be defined as $\mathcal{H} = \{h : X_P^2 \rightarrow \mathbb{R} \mid \exists w \in W, h((x, x')) = \text{sign}(f_w(x) - f_w(x'))\}$ where $\text{sign}(t) = 1$ if $t > 0$ and $\text{sign}(t) = -1$ otherwise, and thus can be identified to the set of admissible weights W . In this set, let w^* denote the weight vector that best represents the DM's preferences, i.e., $w^* = \arg \min_{w \in W} \ell(f_w)$. The supervised learning sub-tasks of the DMH algorithm involve minimizations over W of the empirical error on preference examples $S = (z^i, y^i)_{i=1}^m$ defined as $\ell_S(f_w) = \frac{1}{|S|} \sum_{(x^i, x^{i'}, y^i) \in S} \mathbb{1}_{\{\text{sign}(f_w(x^i) - f_w(x^{i'})) \neq y^i\}}$. However, the 0-1 loss $\ell_S(f_w)$ is non-convex and discontinuous and its optimization on a set W is known to be NP-hard in the case of noisy answers,

even for linear models f_w [12]. To bypass this issue, we use a randomly generated finite approximation $W_0 \subseteq W$ of W and solve the optimization tasks by exhaustive search.

In the next subsection, we introduce an algorithm designed to achieve a twofold objective: on the one hand, finding a near-optimal solution within X_P , and on the other hand, assessing parameter w to have a predictive model f_w of DM’s preferences. To this end, we propose a tailoring of the DMH algorithm to the aforementioned preference learning setting, extended with a mechanism to control the regret of the recommended solutions.

3.2 A Disagreement-based Active Preference Learning Algorithm

Let us introduce Algorithm 1 that takes as input a stream of pairs of alternatives $z^k = (x^k, x'^k)$ drawn randomly and uniformly from the set of alternatives X_P and an initial set of admissible weight vectors W_0 , the definition domain of weight w used in function f_w . This algorithm sequentially proceeds the pairs of alternatives and at iteration k , only when necessary, asks for the DM to provide an answer y^k to the pairwise comparison query $x^k \succ x'^k$?. Since the DM possibly provides answers in contradiction with the preferences induced by f_{w^*} , the obtained preference examples are not exploited as hard constraints to reduce W_0 , but used to construct a growing learning dataset of preference examples T_k .

Algorithm 1:

Inputs: W_0, α, ρ

- 1 draw a pair (x^0, x'^0) uniformly in X_P^2
- 2 $y^0 \leftarrow$ answer to the query “ $x^0 \succ x'^0$?”
- 3 $T_0 \leftarrow \{(x^0, x'^0, y^0)\}, \text{MR}_0 \leftarrow 1, k \leftarrow 1$
- 4 **while** $\text{MR}_{k-1}/\text{MR}_0 > \rho$ **do**
- 5 draw a pair (x^k, x'^k) uniformly in X_P^2
- 6 $W_k^+, W_k^- \leftarrow \{w \in W_{k-1} \mid f_w(x^k) > f_w(x'^k)\}, \{w \in W_{k-1} \mid f_w(x^k) \leq f_w(x'^k)\}$
- 7 **if** $W_k^+ \neq \emptyset$ **and** $W_k^- \neq \emptyset$ **then**
- 8 $w^+, w^- \leftarrow \arg \min_{w \in W_k^+} \ell_{T_{k-1}}(f_w), \arg \min_{w \in W_k^-} \ell_{T_{k-1}}(f_w)$
- 9 $\beta_k \leftarrow \sqrt{(4/k) \ln(8(k^2 + k) |W_0|^2 / \delta)}$
- 10 $\Delta_k \leftarrow \beta_k^2 + \beta_k (\sqrt{\ell_{T_{k-1}}(f_{w^+})} + \sqrt{\ell_{T_{k-1}}(f_{w^-})})$
- 11 **end**
- 12 **if** $W_k^- \neq \emptyset$ **and** $(W_k^+ = \emptyset \text{ or } \ell_{T_{k-1}}(f_{w^+}) - \ell_{T_{k-1}}(f_{w^-}) > \alpha \Delta_k)$ **then**
- 13 $W_k \leftarrow W_k^-, T_k \leftarrow T_{k-1}$
- 14 **else if** $W_k^+ \neq \emptyset$ **and** $(W_k^- = \emptyset \text{ or } \ell_{T_{k-1}}(f_{w^-}) - \ell_{T_{k-1}}(f_{w^+}) > \alpha \Delta_k)$ **then**
- 15 $W_k \leftarrow W_k^+, T_k \leftarrow T_{k-1}$
- 16 **else**
- 17 $y^k \leftarrow$ answer to the query “ $x^k \succ x'^k$?”
- 18 $W_k \leftarrow W_{k-1}, T_k \leftarrow \{(x^k, x'^k, y^k)\}$
- 19 **end**
- 20 $\hat{w}_k \leftarrow \arg \min_{w \in W_k} \ell_{T_k}(f_w)$
- 21 $\hat{x}_k, \text{MR}_k \leftarrow \arg \max_{x \in X_P} f_{\hat{w}_k}(x), \max_{x' \in X_P} \max_{w \in W_k} (f_w(x') - f_w(\hat{x}_k))$
- 22 $k \leftarrow k + 1$
- 23 **end**

Outputs: $\hat{w}_{k-1}, \hat{x}_{k-1}, \text{MR}_{k-1}$

More precisely, at iteration k , to determine whether query “ $x^k \succ x'^k$?” is worth asking (i.e., whether the answer y^k is likely to provide new information), the algorithm assesses the level of disagreement of the current set of admissible models W_{k-1} on the pair (x^k, x'^k) . To this end, W_{k-1} is partitioned into $W_k^+ = \{w \in W_{k-1} \mid f_w(x^k) > f_w(x'^k)\}$ (set of weights that verify $x^k \succ x'^k$) and $W_k^- = \{w \in W_{k-1} \mid f_w(x^k) \leq f_w(x'^k)\}$ (set of weights that verify $x^k \preceq x'^k$) and the minimal empirical errors on the learning database T_{k-1} are computed on both set W_k^+ and W_k^- , and compared. If w^+, w^- respectively denote the best elements in W_k^+ and W_k^- , i.e., $w^+ = \arg \min_{w \in W_k^+} \ell_{T_{k-1}}(f_w)$ and $w^- = \arg \min_{w \in W_k^-} \ell_{T_{k-1}}(f_w)$, this amounts to assessing the gap $\ell_{T_{k-1}}(f_{w^+}) - \ell_{T_{k-1}}(f_{w^-})$.

If the magnitude of the error difference $|\ell_{T_{k-1}}(f_{w^+}) - \ell_{T_{k-1}}(f_{w^-})|$ does not exceed a certain threshold Δ_k , the elements of W_{k-1} somehow “disagree” on whether $x^k \succ x'^k$ or $x^k \preceq x'^k$. Indeed, the weight vectors verifying

$x^k \succ x'^k (W_k^+)$ and the weight vectors verifying $x^k \succsim x'^k (W_k^-)$ are attached to similar minimal errors on the learning database T_{k-1} . Therefore, the answer y^k is likely to provide new information, and thus the query “ $x^k \succ x'^k$?” is asked to the DM. Then, the answer y^k is stored as a new preference example (x^k, x'^k, y^k) in the learning database, i.e., $T_k = T_{k-1} \cup \{(x^k, x'^k, y^k)\}$.

However, if $|\ell_{T_{k-1}}(f_{w^+}) - \ell_{T_{k-1}}(f_{w^-})| > \Delta_k$, the elements of W_{k-1} somehow “agree” on whether $x^k \succ x'^k$ or $x^k \succsim x'^k$, since one of the two sets W_k^+, W_k^- yields significantly higher errors on the learning database T_{k-1} than the other, and thus is not likely to contain the best predictor w^* . Then, the query is not worth asking and the algorithm exploits the agreement of W_{k-1} on the instance (x^k, x'^k) to reduce, with high confidence, the set of admissible models by adding the preference as a hard constraint, i.e., $W_k = W_k^+$, if the constraint $x^k \succ x'^k$ yields the smallest error, and $W_k = W_k^-$ otherwise. In this case, no example is added to the learning database, i.e., $T_k = T_{k-1}$. Note that cases $W_k^+ = \emptyset$ or $W_k^- = \emptyset$ are omitted here for clarity but are included in Algorithm 1.

This cautious reduction of the set of admissible weights makes it possible to incrementally control the remaining level of uncertainty on the DM’s best alternative in X_P , as in incremental preference elicitation methods based on the notion of maximum regret [20, 5, 3]. Indeed, at the end of iteration k , the learned model is $\hat{w}_k = \arg \min_{w \in W_k} \ell_{T_k}(f_w)$ and naturally the recommended solution is $\hat{x}_k = \arg \max_{x \in X_P} f_{\hat{w}_k}(x)$. Then, we propose to assess the remaining level of uncertainty on the DM’s best alternative by computing the maximum regret attached to the recommendation of \hat{x}_k knowing that the current set of admissible weights is W_k , i.e.,: $\text{MR}_k = \max_{x \in X_P} \max_{w \in W_k} \{f_w(x) - f_w(\hat{x}_k)\}$. Once MR_k is sufficiently reduced (ratio MR_k/MR_1 below a configurable threshold $\rho \in [0, 1)$), the algorithm stops and outputs a recommended solution \hat{x}_k .

We now establish Proposition 1 showing how threshold values Δ_k must be set to make sure that $W_k, k \geq 1$ contain with high probability weight vector w^* that best fits to DM’s preferences. This result relies on the known guarantee attached to the DMH algorithm [10] and shows that MR_k upper bounds the real regret $\max_{x \in X_P} \{f_{w^*}(x) - f_{w^*}(\hat{x}_k)\}$, with high probability.

Proposition 1 For $\delta > 0$, $\beta_k = \sqrt{(4/k) \ln(8(k^2 + k) |W_0|^2 / \delta)}$, and $\Delta_k = \beta_k^2 + \beta_k (\sqrt{\ell_{T_{k-1}}(f_{w^+})} + \sqrt{\ell_{T_{k-1}}(f_{w^-})})$, with probability at least $1 - \delta$, we have $w^* \in W_k$ and MR_k upper bounds the real regret, i.e., $\max_{x \in X_P} \{f_{w^*}(x) - f_{w^*}(\hat{x}_k)\} \leq \text{MR}_k$, for any $k \geq 1$, in Algorithm 1.

Algorithm 1 is a specification of the DMH algorithm [10] (with a modified stopping criterion) for the input space $\mathcal{Z} = X_P^2$ associated to the uniform distribution, and the hypothesis class $\mathcal{H} = \{h : X_P^2 \rightarrow \mathbb{R} | \exists w \in W_0, h((x, x')) = \text{sign}(f_w(x) - f_w(x'))\}$. Let \mathcal{H}_k denote the hypothesis class at iteration k , i.e., $\mathcal{H}_k = \{h : X_P^2 \rightarrow \mathbb{R} | \exists w \in W_k, h((x, x')) = \text{sign}(f_w(x) - f_w(x'))\}$. Then, for $\delta > 0$ and $\gamma_k = \sqrt{(4/k) \ln(8(k^2 + k) \mathcal{S}(\mathcal{H}, 2k)^2 / \delta)}$ where $\mathcal{S}(\mathcal{H}, 2k)$ is the k -th shatter coefficient of \mathcal{H} , with probability at least $1 - \delta$, and all $(h, h') \in \mathcal{H}_k \times \mathcal{H}_k$, we have for any $k \geq 1$ ([10], Corollary 1):

$$\ell_{T_{k-1}}(h) - \ell_{T_{k-1}}(h') \leq \ell(h) - \ell(h') + \gamma_k^2 + \gamma_k (\sqrt{\ell_{T_{k-1}}(h)} + \sqrt{\ell_{T_{k-1}}(h')})$$

Recall that the k -th shatter coefficient is defined as the maximal number of ways \mathcal{H} can classify k input points. More formally, $\mathcal{S}(\mathcal{H}, k)$ is the maximal number of different target vector $(h(z_1), \dots, h(z_k))$ that can be associated to k input points (z_1, \dots, z_k) by \mathcal{H} , i.e., $\mathcal{S}(\mathcal{H}, k) = \max_{z \in \mathcal{Z}} |\{(h(z_1), \dots, h(z_k)) | h \in \mathcal{H}\}|$. Then, by definition of the shatter coefficient, when \mathcal{H} is finite, i.e., $|\mathcal{H}| < \infty$, for any $k \geq 1$, $\mathcal{S}(\mathcal{H}, k) \leq |\mathcal{H}|$. Here, $|\mathcal{H}| \leq |W_0| < \infty$ and thus $\mathcal{S}(\mathcal{H}, k) \leq |W_0|$ for any $k \geq 1$. Therefore $\beta_k = \sqrt{(4/k) \ln(8(k^2 + k) |W_0|^2 / \delta)} \geq \gamma_k$, and thus we have, with probability at least $1 - \delta$, for all $(h, h') \in \mathcal{H}_k \times \mathcal{H}_k$ and all $k \geq 1$:

$$\ell_{T_{k-1}}(h) - \ell_{T_{k-1}}(h') \leq \ell(h) - \ell(h') + \beta_k^2 + \beta_k (\sqrt{\ell_{T_{k-1}}(h)} + \sqrt{\ell_{T_{k-1}}(h')})$$

Therefore, for $\Delta_k = \beta_k^2 + \beta_k (\sqrt{\ell_{T_{k-1}}(f_{w^+})} + \sqrt{\ell_{T_{k-1}}(f_{w^-})})$, with probability at least $1 - \delta$, we have $w^* \in W_k$, for any $k \geq 1$ ([10], Lemma 3). Thus, with probability at least $1 - \delta$, $\max_{x \in X_P} \{f_{w^*}(x) - f_{w^*}(\hat{x}_k)\} \leq \max_{w \in W_k} \max_{x \in X_P} \{f_w(x) - f_w(\hat{x}_k)\} = \text{MR}_k$. \square

In the next Section dedicated to numerical tests, we will see that Δ_k is an over-cautious threshold, that may induce many preference queries. In practice, a more aggressive threshold $\alpha \Delta_k$ where $\alpha \in [0, 1]$ can be used to reduce the number of preference queries without too much sacrificing the recommendation quality. Parameter α can be set using cross-validation on a small test set of preference examples $\{(x^i, x'^i, y^i)\}$. On the other side, parameter δ is set to 0.95.

3.3 Illustration on a Toy Example

To illustrate the benefit of Algorithm 1 for preference elicitation with noisy answers, we exploit the easy-to-grasp toy case of Example 1 on the choice set $X = \{a^0, \dots, a^q\}$. The first pair under consideration is $(x^0, x'^0) = (a^r, a^t)$

Table 1: Comparison of Algorithm 1 and CSS over 100 simulations.

	$p = 0.1$		$p = 0.3$	
	Number of query	Rec. accuracy	Number of query	Rec. accuracy
CSS	1.0	89%	1.0	74%
Algorithm 1	3.5	99%	14.0	90%

with $r > t$. Since the first DM's preference statement is $a^r \succ a^t$, Algorithm 1 starts with the initial learning database $T_0 = \{(x^0, x'^0, 1)\}$. Then, examples of pairs $(x^k, x'^k) = (a^{r_k}, a^{t_k})$ are repeatedly drawn uniformly from $X_P^2 = X^2$ (where $r_k > t_k$) and Algorithm 1 either asks the query $x^k \succ x'^k$? or if confident enough to predict the DM's answer, does not ask the query and reduces the current set of admissible weights accordingly (either $(\frac{1}{2}; 1]$ or $[0; \frac{1}{2}]$). In this case, the recommended alternative $\hat{x}_k = \arg \max_{x \in X_P} f_{\hat{w}_k}(x)$ with $\hat{w}_k = \arg \min_{w \in W_k} \ell_{T_k}(f_w)$ is necessarily a^q if $W_k = (\frac{1}{2}; 1]$ or a^0 if $W_k = [0; \frac{1}{2}]$. Finally, the recommended alternative being associated with a null maximum regret (for instance if $W_k = (\frac{1}{2}; 1]$, $\max_{r \in \{0, \dots, q\}} \max_{w \in (1/2, 1]} (f_w(a^r) - f_w(a^q)) = \max_{r \in \{0, \dots, q\}} \max_{w \in (1/2, 1]} \{(r - q)(2w - 1)\} = 0$), the algorithm stops. Also, at each iteration k , one can easily see that $\ell_{T_{k-1}}(f_{w^+}) = \frac{1}{|T_{k-1}|} \sum_{i=0}^{k-1} \mathbb{1}_{\{y^i=1\}}$ and $\ell_{T_{k-1}}(f_{w^-}) = \frac{1}{|T_{k-1}|} \sum_{i=0}^{k-1} \mathbb{1}_{\{y^i=-1\}}$ and thus $\ell_{T_{k-1}}(f_{w^+})$ and $\ell_{T_{k-1}}(f_{w^-})$ are respectively the frequencies of occurrences of preferences of type $x^i \succ x'^i$ and of type $x^i \preceq x'^i$ in the sequence of the past DM's preference statements. Thus, Algorithm 1 recommends a solution when one of the two frequencies becomes significantly higher than the other, i.e., with a difference higher than $\alpha \Delta_k$ which, according to Proposition 1, is decreasing with k . This choice process is more robust than taking for granted the very first DM's preference statement, as illustrated below.

In Table 1, we present numerical results obtained for the described toy case with $q = 20$. To assess the benefit of Algorithm 1 in the context of noisy answers we introduce a random noise in the simulated DM's answers which swaps the answers with probability $p = 0.1$ and $p = 0.3$. For the two noise levels, we compare Algorithm 1 to an incremental preference elicitation method based on the minimization of the maximal regret and the current solution strategy (CSS) [20, 5, 3]. CSS consists of asking to compare solution $x \in X_P$ that minimizes the maximum regret $\text{MR}(x) = \max_{x' \in X_P} \max_{w \in W} \{f_w(x') - f_w(x)\}$ and its best opponent $x' = \arg \max_{x' \in X_P} \max_{w \in W} \{f_w(x') - f_w(x)\}$. Then the set of admissible models W is reduced using the obtained preference as a hard constraint, before iterating on this new set of admissible models. Parameters α and δ of Algorithm 1 are respectively set to $\alpha = 0.05$ and $\rho = 0$. In Table 1, we compare the results of both methods on 500 simulations in terms of number of queries and accuracy of the recommendation (number of simulations where the recommendation was correct). Looking at the results, we can see that CSS, while always terminating after one query, suffers from noisy answers and does not recommend the optimal alternative in only 89% of the time for the low noise level, and 74% of the time for the high noise level. On the contrary, Algorithm 1 recommends the optimal alternative in nearly 100% of the time for the low noise level while asking only 3.5 questions on average, and in 90% of the time for the high noise level with about 14 questions on average. Further tests are conducted in the next section.

4 Numerical Tests

In this section, we present the results of numerical tests performed on synthetic preference data. We test the ability of Algorithm 1 to provide accurate recommendations while receiving noisy answers from the DM, and when possible, we compare those results to CSS. The tests are conducted with f_w taken as the weighted sum, the 2-additive Choquet integral and the Chebyshev distance. We consider random finite sets of admissible models $W_0 \subseteq W$ of size $|W_0| = 5000$ for each experiment. For the weighted sum and the Chebyshev distance, W_0 is obtained by uniform sampling of the simplex $W = \{w \in [0, 1]^n \mid \sum_{i=1}^n w_i = 1\}$. For the 2-additive Choquet integrals, the set of admissible weights is the set of capacities $W = \{w : 2^N \rightarrow \mathbb{R} \mid A \subseteq N, w(A) \leq w(B), w(\emptyset) = 0, w(N) = 1\}$, restricted to 2-additive capacities. The set of 2-additive capacities is a polyhedron admitting a polynomial number of extreme points [14] (Theorem 2.65), namely, the *unanimity games* defined for any $i \in \llbracket 1, \frac{n(n+1)}{2} \rrbracket$ by $v_i(S) = 1$ if $Y_i \subseteq S$ and 0 otherwise, where $Y_i \subseteq N$ is any nonempty subset of size at most 2, and the *conjugates of unanimity games* defined for any $i \in \llbracket \frac{n(n+1)}{2} + 1, n^2 \rrbracket$, by $v_i(S) = 1$ if $Y_i \cap S \neq \emptyset$ and 0 otherwise, for $i \in \llbracket \frac{n(n+1)}{2} + 1, n^2 \rrbracket$. Hence, any 2-additive capacity w can be generated by a convex combination $w = \sum_{i=1}^q \beta_i v_i$ with $q = n^2$, $\beta_i \in [0, 1]$, $\sum_{i=1}^q \beta_i = 1$. Thus, we obtain samples W_0 of the set of 2-additive capacities by uniform sampling of the simplex $\{\beta \in [0, 1]^q \mid \sum_{i=1}^q \beta_i = 1\}$.

For all the experiments, the set X_P of Pareto optimal solutions is randomly drawn within \mathcal{X} as follows: vectors μ of size $n - 1$ are uniformly drawn in $[0, 1]^{n-1}$, then performance vectors $x \in \mathcal{X}$ are obtained by setting $x_i = \mu_{(i)} - \mu_{(i-1)}$ for $i = 1, \dots, n$, where $\mu_{(0)} = 0$ and $\mu_{(p)} = 1$. To avoid that all generated alternatives share the same hyperplane

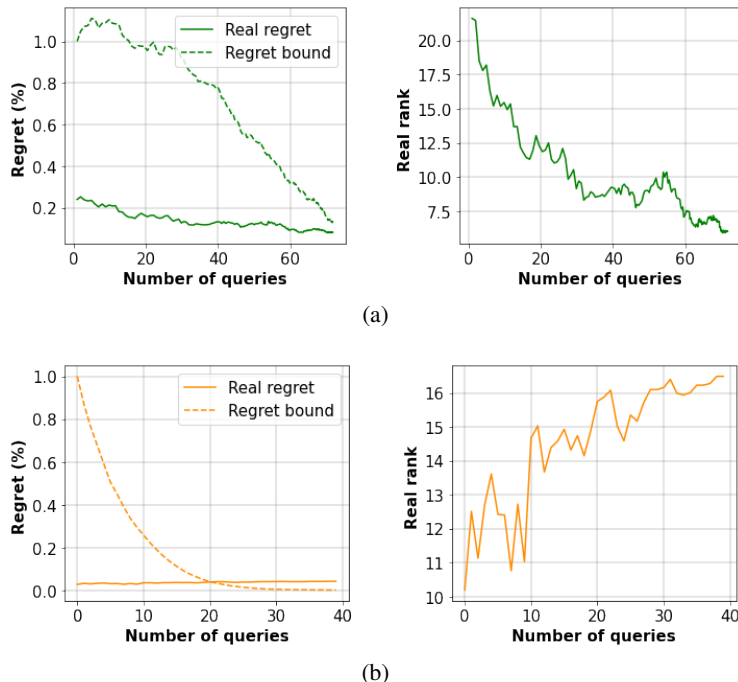


Figure 2: Real regret and rank w.r.t. query number for Algorithm 1 (a) and CSS(b).

$\sum_i x_i = 1$ in the utility space, the square root function is applied on all components x_i of each performance vector x . Also, the DM's answers are simulated according to a ground truth model $f_{w_{gt}}$ with a random weight vector w_{gt} generated in the same way as the elements of W_0 . The answers are disturbed with random noises ϵ such that $y^k = \text{sign}(f_{w_{gt}}(x) - f_{w_{gt}}(x') + \epsilon_k)$ and ϵ_k is uniformly distributed within $[-\sigma, \sigma]$ with noise level $\sigma > 0$, for any k .

In the first experiment, we compare the noise tolerance of Algorithm 1 and CSS when f_w is the Choquet integral, $n = 5$, and $X_P = 100$. Parameter α of Algorithm 1 is set to $\alpha = 1.7 \times 10^{-2}$. Figures 2a and 2b respectively show, for Algorithm 1 and CSS, the average regret (left) and rank in the DM's hidden ranking (right) of the recommended solution over 100 simulations w.r.t. the number of query. More precisely, 2a(left) (resp. 2b(left)) shows the real regret of the recommended solution in percentage w.r.t. the initial level of uncertainty MR_1 (resp. the initial min-max regret $\min_{x \in X_P} \max_{x' \in X_P} \max_{w \in W} \{f_w(x') - f_w(x)\}$) along with the upper bound of the real regret MR_k (resp. the min-max regret of the reduced set of admissible models after k queries). In Figure 2b, we observe that while the min-max regret quickly reduces with CSS, it does not induce a reduction of the real regret and yields recommended solutions with increasing real ranks. On the contrary, in Figure 2a, we observe that while decreasing more slowly, the bound MR_k of Algorithm 1 decreases accordingly with the real regret and rank of the recommended solution \hat{x}_k . After 30 queries, the real rank of the recommended solution is about 8 for Algorithm 1 and 16 for CSS.

In the second experiment, we show the different tradeoffs between quality of the recommendation and number of asked queries that can be achieved with Algorithm 1 by varying the α parameter, which controls the threshold value $\alpha \Delta_k$. The tests are conducted for the Chebyshev distance for $n = 10$, $X_P = 100$ and α varying in a uniform grid within $[5 \times 10^{-3}, 5 \times 10^{-2}]$. The results are averaged over 100 simulations, and for this experiment, the DM's answers are disturbed with a random noise which swaps the answers with probability $p = 0.1$ and $p = 0.2$. For both noise levels, Figure 3a (left) represents the average real regret of the recommended solutions (in percentage w.r.t. MR_1) versus the average number of asked queries and Figure 3a (right) shows the average real rank of the recommended solution, again versus the average number of queries. For all figures, the higher the α value, the higher the caution level of Algorithm 1, and thus the higher the number of asked queries. For $p = 0.1$ (red), asking 7 queries yields a real regret of 20% in average with an average real rank equal to 21 and asking 50 queries reduces the real regret to 10% and the average rank to 8. When the noise level increases, the performances weaken. For instance, for $p = 0.2$, 7 questions yield an average real rank equal to 26.

In the third experiment, we compare Algorithm 1 to another non-Bayesian active learning method recently proposed for linear models [17]. This method also exploits the idea of minimizing the 0-1 loss error on the set of admissible models W instead of irreversibly reducing the set of admissible models such as in CSS. However, while being effective

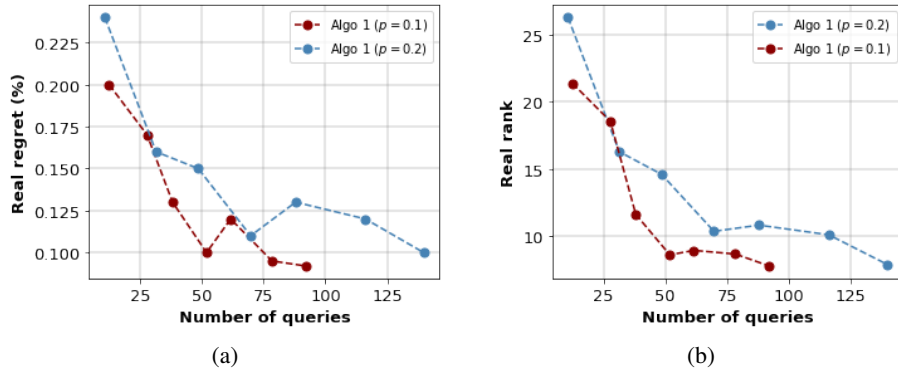


Figure 3: Real regret (a) and rank (b) w.r.t. query number (Chebyshev distance).

at solving choice problems with small number of queries, the used querying strategy focuses only on the most plausible best element of X_P , and thus, the learned model is further from the hidden model w_{gt} and shows lower generalization performances on \mathcal{X} than Algorithm 1. This can be seen in Table 2 where both methods are compared in terms of query number and real rank of the recommended solutions; we also give the average absolute distance to w_{gt} of the learned preference model and the test accuracy defined as the percentage of preference inversion on a test set of pairwise comparison in \mathcal{X} . The tests are conducted with f_w taken as the weighted sum, $n = 10$, $\sigma = 0.05$, $|X_P| = 1000$ and for Algorithm 1 the parameter are set to $\alpha, \rho = (2 \times 10^{-10}, 0.5)$ which allow yielding similar query numbers for both methods. We observe that while yielding similar results in terms of query number and real rank, Algorithm 1 better recovers preference model w_{gt} and achieves a higher test accuracy. Computation times are comparable for both methods (3.35 sec. for Algorithm 1 and 1.33 sec for [Pourkhajouei, 23] on average).

Table 2: Comparison with [Pourkhajouei, 23] over 100 simulations.

	Query number	Real rank	Distance to w_{gt}	Test accuracy
Algo 1	68.0 ± 22.7	$40.1/1000 \pm 59.8$	0.03 ± 0.01	$88.5\% \pm 2.5\%$
[Pourkhajouei, 23]	60.9 ± 17.3	$44.75/1000 \pm 81.13$	0.07 ± 0.02	$78.0\% \pm 5.0\%$

5 Conclusion

We have presented a new approach for determining an optimal solution in a given set, by actively learning the parameters of an aggregation function describing the DM’s preferences. This approach is a cautious version of the standard CSS based on the minimax regret criterion that progressively reduces the set of admissible model parameters, until a zero-regret (or near-zero-regret) solution appears as a necessary winner. In our view, our approach offers three significant advantages.

Firstly, it is more error-tolerant, since the DM’s responses are not systematically interpreted as hard constraints on the parameter space. The numerical tests carried out in Section 4 clearly demonstrated the gain in robustness in the face of noisy responses. The second advantage is that, beyond the identification of an optimal choice, the method provides a learned model that can be used to explain decisions and make choices on new instances. Finally, it does not require the scalarizing function to be linear in its parameters and thus applies to a wider class of aggregators, including the weighted Chebyshev norm, or the Sugeno integral that is generally not learned by regret minimization.

Algorithm 1 also brings some advantages compared to recently proposed approaches for preference learning with noisy DM’s answers, whether Bayesian [6, 4] or non-Bayesian [17]. On the one hand, being non-Bayesian, the proposed approach does not require knowledge of a prior distribution on the model parameters, a strong assumption often necessary to initiate Bayesian learning. On the other hand, concerning non-Bayesian approaches, the numerical tests presented at the end of Section 4 show that Algorithm 1, while exhibiting comparable performance to recent alternative proposals [17] in terms of robustness to noisy responses, achieves significantly better generalization performance and thus is likely to make better decision on new instances of choice problems.

Acknowledgments

This work is supported by the ANR project ANR-20-CE23-0018 THEMIS of the French National Research Agency.

References

- [1] Adam, L., Destercke, S.: Handling inconsistency in (numerical) preferences using possibility theory. *Information Fusion* **103**, 102089 (2024)
- [2] Balcan, M.F., Beygelzimer, A., Langford, J.: Agnostic active learning. In: *Proceedings of the 23rd International Conference on Machine Learning*. pp. 65–72 (2006)
- [3] Benabbou, N., Perny, P., Viappiani, P.: Incremental elicitation of Choquet capacities for multicriteria choice, ranking and sorting problems. *Artificial Intelligence* **246**, 152–180 (2017)
- [4] Bourdache, N., Perny, P., Spanjaard, O.: Incremental elicitation of rank-dependent aggregation functions based on Bayesian linear regression. In: *proceedings of IJCAI-19*. pp. 2023–2029 (2019)
- [5] Boutilier, C., Patrascu, R., Poupart, P., Schuurmans, D.: Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence* **170**(8-9), 686–713 (2006)
- [6] Chajewska, U., Koller, D., Parr, R.: Making rational decisions using adaptive utility elicitation. In: *AAAI/IAAI*. pp. 363–369 (2000)
- [7] Chateauneuf, A., Jaffray, J.Y.: Some characterizations of lower probabilities and other monotone capacities through the use of Möbius inversion. *Mathematical Social Sciences* **17**(3), 263–283 (1989)
- [8] Cohn, D., Atlas, L., Ladner, R.: Improving generalization with active learning. *Machine Learning* **15**, 201–221 (1994)
- [9] Dasgupta, S.: Two faces of active learning. *Theoretical Computer Science* **412**(19), 1767–1781 (2011)
- [10] Dasgupta, S., Hsu, D.J., Monteleoni, C.: A general agnostic active learning algorithm. *Advances in Neural Information Processing Systems* **20** (2007)
- [11] Domshlak, C., Joachims, T.: Unstructuring user preferences: Efficient non-parametric utility revelation. *arXiv preprint arXiv:1207.1390* (2012)
- [12] Feldman, V., Guruswami, V., Raghavendra, P., Wu, Y.: Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing* **41**(6), 1558–1590 (2012)
- [13] Grabisch, M., Marichal, J.L., Mesiar, R., Pap, E.: *Aggregation functions*, vol. 127. Cambridge University Press (2009)
- [14] Grabisch, M., et al.: *Set functions, games and capacities in decision making*, vol. 46. Springer (2016)
- [15] Hanneke, S., et al.: Theory of disagreement-based active learning. *Foundations and Trends® in Machine Learning* **7**(2-3), 131–309 (2014)
- [16] Herin, M., Perny, P., Sokolovska, N.: Learning preference models with sparse interactions of criteria. In: *Proc. of IJCAI* (2023)
- [17] Pourkhajouei, S., Toffano, F., Viappiani, P., Wilson, N.: An efficient non-Bayesian approach for interactive preference elicitation under noisy preference models. In: *European Conference on Symbolic and Quantitative Approaches with Uncertainty*. pp. 308–321. Springer (2023)
- [18] Steuer, R.E.: *Multiple criteria optimization. Theory, Computation, and Application* (1986)
- [19] Vanderpooten, D., Vincke, P.: Description and analysis of some representative interactive multicriteria procedures. In: *Models and Methods in Multiple Criteria Decision Making*, pp. 1221–1238. Elsevier (1989)
- [20] Wang, T., Boutilier, C.: Incremental utility elicitation with the minimax regret decision criterion. In: *IJCAI*. vol. 3, pp. 309–316 (2003)
- [21] White, C.C., Sage, A.P., Dozono, S.: A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics* (2), 223–229 (1984)
- [22] Wierzbicki, A.P.: On the completeness and constructiveness of parametric characterizations to vector optimization problems. *Operations-Research-Spektrum* **8**(2), 73–87 (1986)